National Center for Computational Sciences Snapshot
The Week of August 13, 2007

<u>Shedding Light on Exploding Stars</u>

Researchers simulate key evidence for an accelerating universe

A team led by Stan Woosley at the University of California–Santa Cruz is using the National Center for Computational Sciences (NCCS) Cray XT4 Jaguar supercomputer to unravel the secrets of the universe's biggest thermonuclear explosions.

Known as Type Ia supernovas, these explosions are an indispensable tool to the scientific community, with importance well beyond what they tell us about the life and death of stars. In 1998, they allowed two teams of astronomers to reach an astonishing conclusion: Not only is the universe still expanding, but that expansion is accelerating.

Dubbed the "Breakthrough of the Year for 1998" by *Science* magazine, this epiphany has profound implications. If what we saw of the universe were all there is, the primary force acting on it would be the gravitational pull of stars, galaxies, and other matter, which should be slowing it down. If, however, the universe is accelerating, there must be some other force—dubbed "dark energy"—that pervades empty space and works against gravity.

Type Ia supernovas are, therefore, being used to answer some profound questions. But to fully understand what they are telling us, scientists need to resolve some critical questions about the supernovas themselves. The most important may focus on the correlation between a supernova's light curve and its brightness. Simply put, does this relationship hold for all Type Ia supernovas? Astrophysicists discovered the relationship by observing nearby supernovas whose distance can be verified in other ways. They assume that the relationship holds for more distant explosions as well, an assumption that allows them to gauge the distance to events billions of light-years away. They would love to see proof that their assumption is valid.

Woosley's team will simulate each aspect of the explosion, which begins when a dead star known as a white dwarf pulls hydrogen off a companion star known as a red giant: the long period of churning convection that leads to the supernova's first flames, the expansion of these flames to create an explosion, and the radiation at the end of the event that gives the supernova its brightness. Currently the team is simulating the entire star in its last minute of convection before the explosion, approximating the increasing compression of material as it gets closer to the center of the star. The team looks forward to even more powerful supercomputers such as the as the petascale system planned for the NCCS, which will be able to perform more than 1,000 trillion calculations a second (1 petaflop). With such a system, the team will be able to simulate the convection for longer periods of time without approximating the compression.

The job is exceptionally demanding, with scales ranging from the microscopic to the stellar. The simulations must include the initial flames, which have a thickness in the neighborhood of a ten-thousandth of a centimeter, and the overall event, which can expand to 10 billion kilometers. Without the leadership computing resources of the NCCS, calculations of such size and scale would be impossible.

The Fastest Supercomputers Get More Closet Space

High-Performance Storage System adds capacity and speed

When scientists run simulations on the world's fastest supercomputers, they need to quickly store massive amounts of data. Users at the NCCS speedily squirrel away their data using the High-Performance Storage System (HPSS), designed to meet the big storage demands of big science.

"The goal is to be able to store, at a fast rate, massive amounts of data to keep up with the supercomputers and to be able to handle data coming in from multiple points simultaneously," says Stanley White, senior storage manager at the NCCS, located at Oak Ridge National Laboratory (ORNL) and managed by the U.S. Department of Energy's Office of Science. The center grants supercomputing time to the nation's top scientists.

As the need for storage space grows, the capacity of the HPSS scales from quadrillions of bytes (petabytes) to quintillions of bytes (exabytes) with the addition of disks and tapes. Today the system stores 1.44 petabytes of data—the hard-drive contents of today's typical laptops 9,000 times over. That means scientists using NCCS supercomputers have copious data at their fingertips to address Grand Challenges ranging from elucidating the workings of the human genome to understanding the intricacies of global climate change.

HPSS was created by a consortium of Department of Energy labs—ORNL, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Sandia National Laboratories, and Lawrence Berkeley National Laboratory—and IBM Federal Systems.

"Fifteen years ago, [national] labs realized they needed something of this size," White says. "They recognized Grand Challenge problems were coming up that would require petaflops (calculations executed at a rate of a quadrillion floating operation points per second) of computing power. And they realized those jobs had to have a place to put the data." No product on the market could meet the storage needs of big science. The consortium formed because everybody realized they had to work together to solve the problem. Those developments propagate outside the labs as IBM makes the system commercially available.

When performing at its peak, the HPSS in its current configuration can store a single stream of data as fast as 750 million bits (megabits) per second—equivalent to storing 750 small novels every second. The average speed of a single data stream is 600 megabits per second. (Data storage capacity is expressed as bytes, whereas the speed of data transmission is expressed as bits per second.)

When the storage system is taxed with dealing with a bunch of processes at once, the speed at which data flows through the system—called the aggregate throughput rate—is 3 billion bits (gigabits) per second, White says. With new equipment added this year, the throughput rate during parallel processing will rise to 6 gigabits per second, he says.

Storage capacity is based on the system's architecture. Data is generated in simulations run on the NCCS's Cray XT4 supercomputer, called Jaguar, the world's fastest supercomputer used for open scientific research. The data enters a network populated by "HPSS movers"—individual computers with a 10-gigabit network bandwidth. Movers transfer data over fiber-channel links to 500-gigabyte disks, then on to 500-gigabyte tapes. Since tapes are less expensive than disk drives, the former provide a cost-effective strategy. Working data stays on the computer during processing and is moved to archive storage when runs are complete. Sometimes data is stored to the HPSS during the runs and later retrieved for further processing.

Today, the NCCS stores 8.8 million files from the nation's leading researchers in fusion, combustion, astrophysics, materials science, automobile safety, and more. The climate modelers have the greatest volume of data among the more than two dozen research projects the center supports.

As the Jaguar system grows, White's goal is to assure the HPSS continues to meet the needs of NCCS systems in both throughput and capacity. To that end, HPSS administrators have purchased a high-performance computing communications link called InfiniBand, which will enable an even higher aggregate throughput rate of 9 gigabits per second next year. The system will need testing before it can be employed at the NCCS. Says White: "It will change performance to a very visible degree."

Researchers Get a Hand With Next-Generation Codes

Scaling to Petaflops workshop given by experts from the NCCS, Cray, and AMD

More than 60 computational scientists from around the country gathered at Oak Ridge National Laboratory recently to learn how their codes can take advantage of tens and hundreds of thousands of processors.

Staff from ORNL's NCCS teamed with colleagues from supercomputer maker Cray, Inc., and chipmaker AMD to offer the first NCCS Scaling to Petaflops Workshop, which included both presentations and hands-on assistance. Attendees were shown the architectures for the center's Jaguar supercomputer when it is upgraded to 250 trillion calculations a second (250 teraflops) and for the center's upcoming 1,000 teraflops (1 petaflop) system. They learned about AMD's upcoming quad-core computer chips. And they received detailed instructions on designing a scalable application.

They also got a glimpse of the process from Bronson Messer of the NCCS Scientific Computing Group, who discussed the challenges of scaling code for a supernova simulation, and David Lignell of Sandia National Laboratories, who talked about the process as it applies to combustion simulation.

Representatives from the NCCS Scientific Computing Group and Cray also worked directly with workshop attendees, demonstrating how to characterize application bottlenecks using Cray performance tools and identify alternative methods and algorithms for the codes.

The workshop will help attendees take advantage of future supercomputers, using the growing number of computing cores offered by such machines to develop groundbreaking science. Attendees also agreed that it was valuable, according to Ravindra Uplenchwar of the Stanford Linear Accelerator Center in California. Uplenchwar said he especially appreciated the hands-on help he received.

“I had some expectations about the workshop,” he said, “and I would think they were completely met.”